arXiv:1411.1451v1 [stat.ME] 5 Nov 2014

Modelling extremes using approximate Bayesian Computation

R. Erhardt* and S. A. Sisson†

Abstract

By the nature of their construction, many statistical models for extremes result in likelihood functions that are computationally prohibitive to evaluate. This is consequently problematic for the purposes of likelihood-based inference. With a focus on the Bayesian framework, this chapter examines the use of approximate Bayesian computation (ABC) techniques for the fitting and analysis of statistical models for extremes. After introducing the ideas behind ABC algorithms and methods, we demonstrate their application to extremal models in stereology and spatial extremes.

1 Introduction

Suppose interest is in modelling the extremes of a multivariate random process. A useful example to hold in mind might be measurements of temperature y sampled at locations $x_1, ..., x_D$, where there is dependence among the D locations due to their proximity to one another. The extremal dependence may, in general, differ from the dependence of non-extremes, and so the model should target the extremes only and not allow the bulk of non-extreme data to overwhelm to model fit. Models for extremes are useful when trying to estimate the risks associated with rare but influential events.

*Wake Forest University, U. S. A.
†School of Mathematics and Statistics, University of New South Wales

Call y_{obs} the observed data from a random process $f(y|\theta)$, where θ is the parameter, and call $\mathcal{A}$ an event or set of events of particular concern. The probability of such an event (with notation to remind of the influence from the parameter θ) is

$$p_\theta = Pr(y_{obs} \in A) = \int_{\mathcal{A}} f(y|\theta)dy,$$

and p_θ^{-1} is termed the *return period* of the event $\mathcal{A}$. A return period of twenty years means that the event occurs roughly once every twenty years. Estimating return periods of extremes is crucial to designing systems capable of handling such extremes. Examples include flood walls that can withstand the 100-year flood, infrastructure that can withstand extreme heat or cold, insurance companies that can remain solvent after losses of a particular magnitude, and so forth. Given the sparsity of data on extreme events, information on θ can be minimal, and this can translate into large parameter uncertainty in θ. Coles and Powell (1996) and **Stephenson (earlier chapter in this book)** advocate the Bayesian approach to inference, arguing that expert opinion incorporated through a prior distribution $p(\theta)$ could be of tremendous value given the natural scarcity of observed data extremes. Alternatively, if one chooses a vague prior $p(\theta)$ and obtains a posterior $\pi(\theta|y_{obs})$, then the parameter uncertainty in θ would naturally be incorporated into calculations, such as predictive return levels.

Now, consider models for multivariate extremes. Suppose that $Y_1, ..., Y_n$ are univariate, independent and identically distributed replicates from some distribution function F, and define $M_n = \max(Y_1, ..., Y_n)$ as the maximum of the n random variables. The distribution of M_n can be obtained exactly assuming F is known. In practice, then one could estimate F from all of the data $Y_1, ..., Y_n$ and estimate the distribution of M_n as $P(M_n \leq z) = F^n(z)$, but this approach has two drawbacks. The first is that even minor discrepancies in estimating F result in large discrepancies in F^n, particularly in the tails of F. Put another way, why should a model which fits the bulk of data to also be a good fit in the tails? A second drawback is that in the limit as $n \to \infty$, F^n does not converge to a non-degenerate distribution. Instead,

one may model renormalized maxima $(M_n - b_n)/a_n$ for sequences $a_n > 0$ and b_n. If there exist sequences $a_n > 0$ and b_n such that

$$\lim_{n\to\infty} P\left(\frac{M_n - b_n}{a_n} \leq z\right) \to G(z)$$

for some non-degenerate distribution function G, then G is a member of the *Generalized Extreme Value* (GEV) family, with distribution function

$$G(z) = \exp\left[-\left(1 + \xi\frac{z-\mu}{\sigma}\right)_+^{-1/\xi}\right]. \tag{1}$$

Here $a_+ = \max(a, 0)$, and μ, σ, and ξ are the location, scale, and shape parameters, respectively (Coles et al., 2001; Gnedenko, 1944). The sign of the shape parameter ξ determines that G corresponds to one of the three classical extreme values distributions: $\xi > 0$ is Fréchet with support $z \in [\mu - \sigma/\xi, +\infty)$, $\xi < 0$ is Weibull with support $z \in (-\infty, \mu - \sigma/\xi]$, and $\xi \to 0$ is Gumbel with support $z \in (-\infty, +\infty)$. The Generalized Extreme Value distribution G has the property of max-stability, understood as follows: if $Y_1, ..., Y_n$ are i.i.d. draws from G, then $\max(Y_1, ..., Y_n)$ also has distribution G, meaning

$$G^n(A_n z + B_n) = G(z)$$

for appropriate sequences $A_n > 0$ and B_n. In fact, a distribution is max-stable if and only if it is a member of the GEV family (Leadbetter and Lindgren, 1983). If block maxima are taken over a block size large enough to allow the GEV to be a valid approximation, then if one further increased the block size (from monthly to annual maxima, for example) the GEV model would still hold, with only a change in the three parameters. While these results for the GEV family assume i.i.d. data, this assumption can be relaxed and the limiting distribution still holds so long as certain mixing conditions are satisfied (Leadbetter and Lindgren, 1983).

A useful member of the GEV family is the unit-Fréchet distribution, with distribution

function

$$P(Z \leq z) = \exp\left(-\frac{1}{z}\right).$$

The simplicity of the unit-Fréchet distribution function is helpful when one considers multivariate and ultimately spatial extremes. Any member of the GEV family may be transformed to have unit-Fréchet margins as follows: if Z has a GEV distribution, and a new variable U is defined as

$$U = \left(1 + \xi\frac{Z-\mu}{\sigma}\right)^{1/\xi}, \tag{2}$$

then U has unit-Fréchet margins. This transformation assumes that the parameters are known. If the parameters are unknown, they may first be estimated and then the transformation to U is taken. For extreme values data in a spatial setting, the first step is often to transform data at each location to unit-Fréchet by fitting all marginal distributions. Then the second step is to analyze the spatial dependence among sites once every location has been transformed. Thus, there is no loss of generality when one assumes unit-Fréchet margins.

This approach can be extended to handle multivariate extremes. Let $(X_{i1}, ..., X_{iD})$, $i = 1, ..., n$ be a $D-$dimensional random vector and let $M_n = (M_{n1}, ..., M_{nD})$ be the vector of componentwise maxima, where $M_{nd} = \max(X_{1d}, ..., X_{nd})$ for $d = 1, ..., D$. It is worth noting that M_n will not appear in the data record unless the occurrence times of each element's block maximum happen to coincide. In a spatial context, this vector M_n might refer to the annual maxima of some variable at D locations. A non-degenerate limit for M_n exists if there exist sequences $a_{nd} > 0$ and b_{nd}, $d = 1, ..., D$ such that

$$\lim_{n\to\infty} P\left(\frac{M_{n1} - b_{n1}}{a_{n1}} \leq z_1, \ldots, \frac{M_{nD} - b_{nD}}{a_{nD}} \leq z_D\right) = G(z_1, \ldots, z_D).$$

Then G is a multivariate extreme value distribution (MEVD), and is max-stable if there exist sequences $A_{nd} > 0$, B_{nd}, $d = 1, ..., D$ such that, for any $n > 1$

$$G^n(z_1, \ldots, z_D) = G(A_{n1}z_1 + B_{n1}, \ldots, A_{nD}z_D + B_{nD}).$$

The univariate marginal distributions of a multivariate extreme value distribution are all necessarily GEV distributions. Thus, for each margin one can define a transformation with parameter (μ_d, σ_d, ξ_d) and transform to unit-Fréchet using equation (2). Since all GEV distributions can be transformed into unit-Fréchet, all MEVD can be transformed into multivariate unit-Fréchet, and thus one may assume, without loss of generality, that all MEVD have unit Fréchet margins. This works out because the domain of attraction condition is preserved under monotone transformations of the marginal distributions (Resnick, 1987). For D fixed locations $x_1, \ldots, x_D$, the joint distribution function can be written as

$$P(Z(x_1) \leq z_1, \ldots, Z(x_D) \leq z_D) = \exp\left(-V(z_1, \ldots, z_D)\right) \tag{3}$$

where $V(z_1, ..., z_D)$ is the exponent measure first described by Pickands (1981). This function takes the form

$$V(z) = D \cdot \int_{\Delta_D} \max_{d=1,\ldots,D} \frac{w_d}{z_d} H(\, dw) \tag{4}$$

where $\Delta_D = w \in \mathbb{R}_+^D \mid w_1 + \ldots + w_D = 1$ is the $D-1$ dimensional simplex, and the angular (or spectral) measure H is a probability measure on Δ_D which determines the dependence structure of the random vector. Due to the common marginal distributions, H has moment conditions $\int_{\Delta_D} w_d H(\, w) = 1/D$ for $d = 1, ..., D$. Max-stability implies that for all N,

$$\begin{aligned} P(Z_1 \leq z_1, \ldots, Z_D \leq z_D)^N &= \exp(-N \cdot V(z_1, \ldots, z_D)) \\ &= \exp(-V(z_1/N, \ldots, z_D/N)) \end{aligned}$$

with the final equality following from the homogeneity property of the exponent measure. The measure also satisfies two bounds: if all locations are independent, $V(z_1, ..., z_D) = 1/z_1 + ... + 1/z_D$; if all locations are totally dependent, $V(z_1, \ldots, z_D) = \max(1/z_1, \ldots, 1/z_D)$.

Thus, we always have

$$\max(1/z_1, \ldots, 1/z_D) \leq V(z_1, \ldots, z_D) \leq 1/z_1 + \ldots + 1/z_D.$$

There are two challenges to working with the spectral representation of the joint distribution function shown in equation (3). First, even if a closed form for the exponent measure can be found by solving equation (4), the joint density function undergoes a combinatorial explosion as the dimension D increases. Differentiating $\exp(-V)$ with respect to the values $z_1, ..., z_D$ leads to a rapid growth in terms:

$$\begin{aligned}
\frac{\delta}{\delta z_1} \exp(-V) &= -V_1 \exp(-V) \\
\frac{\delta^2}{\delta z_1 z_2} \exp(-V) &= (V_1 V_2 - V_{12}) \exp(-V) \\
\frac{\delta^3}{\delta z_1 z_2 z_3} \exp(-V) &= (-V_1 V_2 V_3 + V_{12} V_3 + V_{13} V_2 + V_{23} V_1 - V_{123}) \exp(-V) \\
&\ldots
\end{aligned}$$

where V_i is the partial derivative of V with respect to z_i. Thus even if a reasonable choice for V can be found, as the dimension D increases one is left with an unwieldy likelihood function, which may be difficult to maximize. More common, though, is the situation where closed-form expressions for the exponent measure cannot be obtained by solving equation (4).

As a result, the lack of a closed-form likelihood presents a stumbling block for modelling high dimensional extremes data. While there are a number of procedures that may permit some form of statistical inference in this setting (see the other chapters in this Handbook), the remainder of this chapter will demonstrate how approximate Bayesian computing can be implemented as one solution to this problem in the Bayesian framework.

2 A primer on approximate Bayesian computation

Suppose that interest is in performing a standard Bayesian analysis for a parameter $\theta \in \Theta$ which has a prior distribution $p(\theta)$. The model is defined through the likelihood function $L(y|\theta)$, which is assumed to be a candidate for the true data generating process that produced the observed dataset $y_{obs} \in \mathcal{Y}$. The posterior distribution of the model parameter θ, having now observed the data y_{obs} through the likelihood function, is expressed as $\pi(\theta|y_{obs}) \propto L(y_{obs}|\theta)p(\theta)$. In the present setting, we could have e.g. $\theta = (\mu, \sigma, \xi)$ as the parameters of a generalised extreme value distribution. The posterior distribution contains all the information that is needed to perform inference on the model (see e.g. Gelman et al. (2003); O'Hagan and Forster (2004)).

As the posterior distribution is rarely available in closed form, it is common to base subsequent analysis on a Monte Carlo approximation to the posterior (Brooks et al., 2011). In this manner, if $\theta^{(1)}, \ldots, \theta^{(N)} \sim \pi(\theta|y_{obs})$ are N samples drawn from the posterior distribution, then the posterior expectation of some function $a(\theta)$ can be approximated by

$$\mathbb{E}_\pi[a(\theta)] = \int_\Theta a(\theta)\pi(\theta|y_{obs})d\theta \approx \frac{1}{N}\sum_{i=1}^{N} a(\theta^{(i)}),$$

where $\mathbb{E}_\pi$ denotes expectation under π. There are a wide variety of algorithms available to draw samples from the posterior distribution (Doucet et al., 2001; Chen et al., 2000; Brooks et al., 2011). One of the simplest of these is importance sampling, as illustrated in Algorithm 1, which produces weighted samples $(w^{(1)}, \theta^{(1)}), \ldots, (w^{(N)}, \theta^{(N)})$ from $\pi(\theta|y_{obs})$ based on samples $\theta^{(1)}, \ldots, \theta^{(N)}$ from a sampling distribution $g(\theta)$. In this setting, and writing $w(\theta) = \pi(\theta|y_{obs})/g(\theta)$, then

$$\begin{aligned}\mathbb{E}_g[w(\theta)a(\theta)] &= \int_\Theta w(\theta)a(\theta)g(\theta)d\theta = \int_\Theta a(\theta)\pi(\theta|y_{obs})d\theta = \mathbb{E}_\pi[a(\theta)] \\ &\approx \sum_{i=1}^{N} w(\theta^{(i)})a(\theta^{(i)}).\end{aligned}$$

That is, appropriately weighted samples from $g(\theta)$ can be used as samples from $\pi(\theta|y_{obs})$.

Simple importance sampling algorithm

Input:
An observed dataset, y_{obs}.
A desired number of samples $N > 0$.
A sampling distribution $g(\theta)$, with $g(\theta) > 0$ if $p(\theta) > 0$.

Iterate: For $i = 1, \ldots, N$:
1. Sample a parameter vector from sampling distribution $\theta^{(i)} \sim g(\theta)$.
2. Weight each sample $\theta^{(i)}$ by $w^{(i)} \propto \pi(\theta^{(i)}|y_{obs})/g(\theta^{(i)})$.

Output:
A set of N weighted samples $(w^{(1)}, \theta^{(1)}), \ldots, (w^{(N)}, \theta^{(N)})$ drawn from $\pi(\theta|y_{obs})$.

Table 1: A simple importance sampling algorithm, based on a single large sample of size N.

Almost all posterior simulation algorithms need to be able to evaluate the likelihood function $L(y|\theta)$ in order to be correctly implemented. In the importance sampling algorithm (Algorithm 1) this occurs in the weight evaluation $w^{(i)} \propto L(y_{obs}|\theta^{(i)})p(\theta^{(i)})/g(\theta^{(i)})$. In the present setting, the natural construction of many useful statistical models for extremes results in the likelihood function being computationally prohibitive to evaluate (see Section 1). This computational intractability means that for these classes of models, an alternative procedure is needed to sample from the posterior distribution *without* directly evaluating the likelihood function. One class of procedures that has been developed to achieve this is known as approximate Bayesian computation (Beaumont, 2010; Bertorelle et al., 2010; Csilléry et al., 2010; Sisson and Fan, 2011; Marin et al., 2012).

2.1 Approximate Bayesian computation basics

All approximate Bayesian computation (ABC) procedures operate on the following heuristic argument, that was first developed in the population genetics literature (Tavaré et al., 1991; Pritchard et al., 1999). Suppose that we have a parameter vector $\theta^{(i)}$ that is a candidate draw from the posterior distribution $\pi(\theta|y_{obs})$. Further suppose that we can quickly generate an auxiliary dataset from the model, conditional on this candidate parameter vector, so

that $y^{(i)} \sim L(y|\theta^{(i)})$. Now, the argument states that if $y^{(i)}$ and y_{obs} are "close" to each other (in a general sense that will be made more precise below), then it is credible that the parameter vector $\theta^{(i)}$ could also have generated the observed dataset y_{obs}. In which case, the parameter vector $\theta^{(i)}$ should be retained as an approximate sample from $\pi(\theta|y_{obs})$. Conversely, if $y^{(i)}$ and y_{obs} are not "close" to each other, then $\theta^{(i)}$ is unlikely to be able to generate the observed dataset, and so it should be discarded as not being a draw from the posterior. Repeating this procedure will produce samples that are either approximate draws from the posterior $\pi(\theta|y_{obs})$, or exact draws from some as yet unspecified approximation to the posterior distribution. Either way, direct numerical evaluation of the computationally intractable likelihood function has been avoided.

The principles behind the above heuristic method can be made more precise. Suppose that we generate our candidate parameter vectors from the prior $\theta^{(i)} \sim p(\theta)$. (As part of an importance sampling algorithm, the candidate parameter vectors $\theta^{(i)}$ may be generated from $g(\theta)$ and then reweighted.) Then, given this, a dataset is generated from the likelihood. In this way, the pair

$$(\theta^{(i)}, y^{(i)}) \sim L(y|\theta)p(\theta)$$

has been generated from the prior predictive distribution $p(y, \theta) = L(y|\theta)p(\theta)$. In general, the auxiliary and observed dataset can be compared via some distance metric $\|y^{(i)} - y_{obs}\|$, such as Euclidean or Mahalanobis distance. The retention and discarding of "close" and not "close" auxiliary datasets can be mimicked by computing an importance weight $K_h(\|y^{(i)} - y_{obs}\|)$ where $K_h(u) = K(u/h)/h$ is a standard smoothing kernel with scale parameter $h > 0$. For example, if K_h corresponds to the uniform kernel over $(-h, h)$, then $\theta^{(i)}$ will receive the weight 1 if $\|y^{(i)} - y_{obs}\| \leq h$ (i.e. it is retained if $y^{(i)}$ and y_{obs} are sufficiently "close") and the weight 0 if $\|y^{(i)} - y_{obs}\| > h$ (i.e. it is rejected). Other choices of kernel K_h will produce continuous weights.

The resulting weighted samples $(w^{(i)}, \theta^{(i)}, y^{(i)})$ are then draws from

$$\pi_h^{ABC}(\theta, y|y_{obs}) \propto K_h(\|y - y_{obs}\|)L(y|\theta)p(\theta) \tag{5}$$

(Reeves and Pettitt, 2005; Wilkinson, 2008). Distribution (5) is the joint posterior distribution of model parameter θ and auxiliary dataset y such that y and y_{obs} are "close" in a specific sense. If we are just interested in the resulting distribution of the parameter vector θ as an approximation of the posterior distribution $\pi(\theta|y_{obs})$, then

$$\pi_h^{ABC}(\theta|y_{obs}) \propto \int_{\mathcal{Y}} K_h(\|y - y_{obs}\|)L(y|\theta)p(\theta)dy. \tag{6}$$

Equation (6) is the ABC approximation to the true posterior distribution $\pi(\theta|y_{obs})$.

The quality of this approximation is determined by the kernel scale parameter $h > 0$. Consider what happens to $\pi_h^{ABC}(\theta|y_{obs})$ as h gets small. We have

$$\begin{aligned}
\lim_{h\to 0} \pi_h^{ABC}(\theta|y_{obs}) &\propto \lim_{h\to 0} \int_{\mathcal{Y}} K_h(\|y - y_{obs}\|)L(y|\theta)p(\theta)dy \\
&= \int_{\mathcal{Y}} \mathbf{1}(y = y_{obs})L(y|\theta)p(\theta)dy \\
&= L(y_{obs}|\theta)p(\theta) \\
&\propto \pi(\theta|y_{obs}),
\end{aligned}$$

where $\mathbf{1}(A) = 1$ if A is true, and 0 otherwise. That is, if we only accept those candidate parameter values $\theta^{(i)}$ that exactly reproduce the observed dataset, then we will exactly recover the true posterior distribution $\pi(\theta|y_{obs})$. In practice however, unless y_{obs} is discrete, this can never occur (and in fact, will be unlikely to occur practically for discrete y_{obs} in all but trivial analyses). Hence h will typically be greater than zero in practice.

In the other extreme

$$\begin{aligned}\lim_{h\to\infty} \pi_h^{ABC}(\theta|y_{obs}) &\propto \lim_{h\to\infty} \int_{\mathcal{Y}} K_h(\|y - y_{obs}\|) L(y|\theta) p(\theta) dy \\ &\propto \int_{\mathcal{Y}} L(y|\theta) p(\theta) dy \\ &= p(\theta),\end{aligned}$$

assuming that $K_h(u) \propto 1$ as $h \to \infty$. That is, sampling from the prior $p(\theta)$ and then not showing any discrimination in favour of auxiliary data closely matching the observed data will simply result in all samples being equally weighted, and the ABC approximation to the posterior distribution being given by the prior $p(\theta)$. It should be clear that increasing fidelity toward reproducing the observed data, as measured by decreasing h, defines a smooth transition from prior $p(\theta)$ to posterior $\pi(\theta|y_{obs})$, and that in practice the actually attainable ABC posterior approximation $\pi_h^{ABC}(\theta|y_{obs})$ lies somewhere between the two. Given that it results in greater closeness to $\pi(\theta|y_{obs})$, for inferential purposes lower h is desirable.

Aside from the approximation to the posterior linked to h, in practice a second level of approximation is commonly introduced in an ABC analysis, and one which can have a greater impact on the quality of the ABC approximation of $\pi(\theta|y_{obs})$ than h. To understand the motivation for this, consider the form of the likelihood component of $\pi_h^{ABC}(\theta|y_{obs})$ (6) where, for simplicity, both y_{obs} and θ are univariate. Using a Taylor expansion and the substitution $u = (y - y_{obs})/h$, then

$$\int_{\mathcal{Y}} K_h(\|y - y_{obs}\|) L(y|\theta) dy \approx L(y_{obs}|\theta) + \frac{1}{2} h^2 L''(y_{obs}|\theta) \int u^2 K(u) du,$$

assuming the usual kernel function properties of $\int K(u)du = 1$, $\int uK(u)du = 0$ and $K(u) = K(-u)$. That is, the ABC simulation procedure is simply performing a form of conditional kernel density estimation, by estimating a smoothed likelihood function, and then using this as part of a regular Bayesian analysis (Blum, 2010).

Under this interpretation, a limitation of the proposed ABC method becomes apparent. Kernel density estimation is well known to suffer from the curse of dimensionality, and is arguably impractical in more than two dimensions. Here, the appropriate dimensionality is in $y - y_{obs}$ (although it is masked by the univariate measure $\|y - y_{obs}\|$) within the kernel function K_h. That is, as the dimension of y_{obs} increases, the performance of the ABC posterior approximation $\pi_h^{ABC}(\theta|y_{obs})$ will rapidly deteriorate (Blum, 2010). In effect, it becomes increasingly unlikely to be able to reproduce the observed dataset y_{obs} by randomly sampling y, as the dimension of y_{obs} increases, thereby forcing the practitioner to increase h for a fixed number of samples N.

A simple solution to this is to reduce the dimensionality of y_{obs} by reducing it to a vector of summary statistics $s_{obs} = S(y_{obs})$ and then performing the ABC procedure as before, but using s_{obs} rather than y_{obs} (Tavaré et al., 1991; Pritchard et al., 1999; Beaumont et al., 2002). In this manner, (5) and (6) become

$$\begin{aligned} \pi_h^{ABC}(\theta, s|s_{obs}) &\propto K_h(\|s - s_{obs}\|)L(s|\theta)p(\theta) \quad (7) \\ \pi_h^{ABC}(\theta|s_{obs}) &\propto \int_{\mathcal{S}} K_h(\|s - s_{obs}\|)L(s|\theta)p(\theta)ds, \quad (8) \end{aligned}$$

where $\mathcal{S} = \{S(y) : y \in \mathcal{Y}\}$ is the image of $\mathcal{Y}$ under S, and where $L(s|\theta)$ corresponds to the likelihood function of the summary statistic, which is also assumed to be computationally intractable. In the case where $S(y) = y$ then (7) and (8) reduce to (5) and (6).

Note that once a decision has been made on the summary statistics, then the most accurate possible ABC posterior approximation is given by $\pi(\theta|s_{obs})$ as

$$\begin{aligned} \lim_{h\to 0} \pi_h^{ABC}(\theta|s_{obs}) &\propto \lim_{h\to 0} \int_{\mathcal{S}} K_h(\|s - s_{obs}\|)L(s|\theta)p(\theta)ds \\ &= \int_{\mathcal{S}} \mathbf{1}(s = s_{obs})L(s|\theta)p(\theta)ds \\ &= L(s_{obs}|\theta)p(\theta) \\ &\propto \pi(\theta|s_{obs}), \end{aligned}$$

in the same manner as before. Consequently it is important that due consideration is given to the summary statistics aspect of the model.

In the case of sufficient statistics, there is no loss of information and so $\pi(\theta|s_{obs}) = \pi(\theta|y_{obs})$. When $S(y)$ is not sufficient, then there is some information loss, and $\pi(\theta|s_{obs})$ will be less precise than $\pi(\theta|y_{obs})$. In turn, the ABC posterior approximation $\pi_h^{ABC}(\theta|s_{obs})$ will be less precise than $\pi(\theta|s_{obs})$. While the reduction in dimension from y_{obs} to s_{obs} will allow for increased precision by permitting a reduced kernel scale parameter h, this must be offset by any loss of information in the construction of the summary statistics (Blum et al., 2013).

ABC importance sampling algorithm

Input:
An observed dataset, y_{obs}.
A desired number of samples $N > 0$.
A sampling distribution $g(\theta)$, with $g(\theta) > 0$ if $p(\theta) > 0$.
A smoothing kernel K_h and scale parameter $h > 0$.
A low-dimensional vector of summary statistics $s = S(y)$.
Compute $s_{obs} = S(y_{obs})$.

Iterate: For $i = 1, \ldots, N$:
1. Sample a parameter vector from sampling distribution $\theta^{(i)} \sim g(\theta)$.
Simulate a dataset from the likelihood given parameter vector $\theta^{(i)}$ as $y^{(i)} \sim L(y|\theta^{(i)})$.
Compute the summary statistics $s^{(i)} = S(y^{(i)})$.
2. Weight each sample $\theta^{(i)}$ by $w^{(i)} \propto K_h(\|s^{(i)} - s_{\text{obs}}\|)p(\theta^{(i)})/g(\theta^{(i)})$.

Output:
A set of $i = 1, \ldots, N$ weighted samples $(w^{(i)}, \theta^{(i)}, s^{(i)})$, drawn from $\pi_h^{ABC}(\theta, s|s_{obs})$ (7).
Or a set of $i = 1, \ldots, N$ weighted samples $(w^{(i)}, \theta^{(i)})$, drawn from $\pi_h^{ABC}(\theta|s_{obs})$ (8).

Table 2: A simple ABC importance sampling algorithm, based on a single large sample of size N.

A precise importance sampling algorithm to generate from $\pi_h^{ABC}(\theta|s_{obs})$ (8) is given by Algorithm 2. To see how this works while avoiding evaluation of the intractable likelihood function $L(s|\theta)$, note that this algorithm targets the joint posterior approximation $\pi_h^{ABC}(\theta, s|s_{obs})$ – marginalising over s so that the draws come from $\pi_h^{ABC}(\theta|s_{obs})$ is achieved by simply discarding the $s^{(i)}$ in the Monte Carlo output. Hence our sampling distribution

must span (θ, s). In the present setting, it is natural to use

$$(\theta^{(i)}, s^{(i)}) \sim L(s|\theta)g(\theta).$$

From which the importance weight becomes

$$\begin{aligned} w^{(i)} &\propto \frac{\pi_h^{ABC}(\theta^{(i)}|s_{obs})}{L(s^{(i)}|\theta^{(i)})g(\theta^{(i)})} \propto \frac{K_h(\|s^{(i)} - s_{obs}\|)L(s^{(i)}|\theta^{(i)})p(\theta^{(i)})}{L(s^{(i)}|\theta^{(i)})g(\theta^{(i)})} \\ &= \frac{K_h(\|s^{(i)} - s_{obs}\|)p(\theta^{(i)})}{g(\theta^{(i)})}, \end{aligned}$$

which is conveniently free of computationally intractable likelihood terms.

2.2 A simple example

As an illustration of ABC methods in a simple setting, consider a standard analysis of univariate block maxima. Figure 1 illustrates the annual maximum daily rainfall values (measured in millimetres) in the years 1951–1999, recorded at Maiquetia International Airport, on the central coast of Venezuela. This dataset is particularly notable for the unusually extreme daily rainfall events in December 1999 that caused the worst environmentally related tragedy in Venezuelan history, and one of the largest historical rainfall-induced debris flows documented in the world (Wieczorek et al., 2001). The resulting human, infrastructure and economic impacts have become known as the *Vargas Tragedy*. These data have previously been analysed by Coles and Pericchi (2003) and Coles et al. (2003).

As these data are univariate annual maxima, they may be approximately modelled by the generalised extreme value (GEV) distribution, which has distribution function shown in (1). From this, and given a prior specification $p(\mu, \sigma, \xi)$, samples can easily be obtained from the posterior distribution $\pi(\mu, \sigma, \xi|y_{obs})$ using standard posterior simulation algorithms, such as Markov chain Monte Carlo (MCMC).

As the GEV is obviously tractable for analysis, this means that a gold standard result

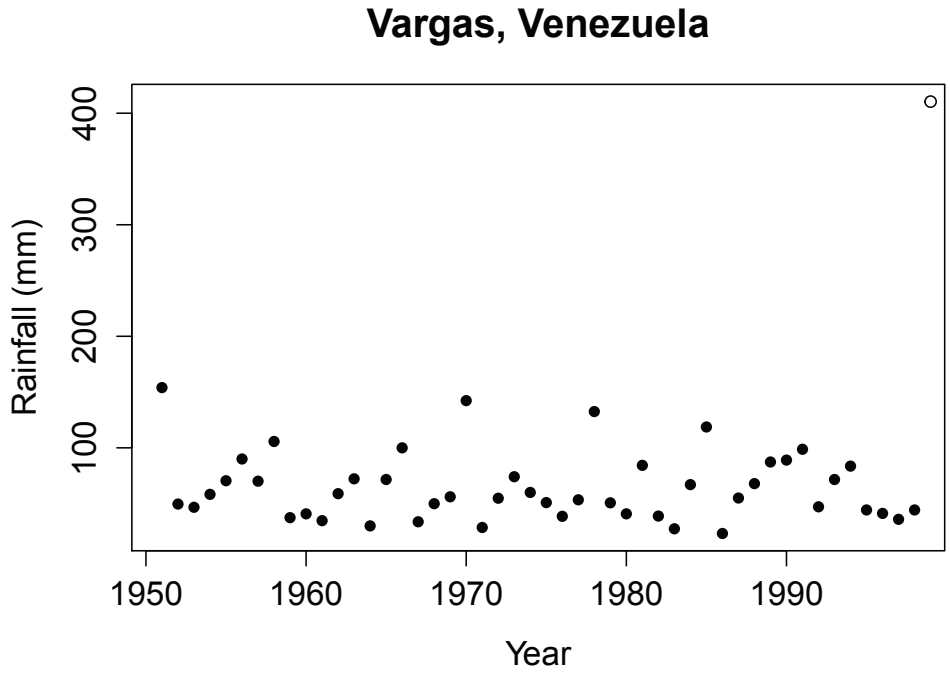

Figure 1: Annual maximum daily rainfall values (mm) 1951–1999 recorded at Maiquetia station on the central coast of Venezuela. The open circle denotes the extreme rainfall observed during the *Vargas Tragedy*, in December 1999.

is available to compare with an ABC posterior approximation. Four specifications of the vector of summary statistics are considered for the dataset $y = (y_1, \ldots, y_n)$:

$$\begin{aligned} s_1(y) &= (y_1, \ldots, y_n)' = y \\ s_2(y) &= (y_{(1)}, \ldots, y_{(n)})' \\ s_3(y) &= (\hat{\mu}_L, \hat{\sigma}_L, \hat{\xi}_L)' \\ s_4(y) &= (\hat{\mu}, \hat{\sigma}, \hat{\xi})'. \end{aligned}$$

Here, $s_1(y) = y$ is the original full dataset, whereas $s_2(y)$ is the vector of n order statistics of y such that $y_{(1)} \leq \ldots \leq y_{(n)}$. The vector $s_3(y)$ consists of the three L-moments estimates (following Hosking (1990)) of each GEV parameter, and $s_4(y)$ is formed from the standard maximum likelihood estimates based on (1). Note that $s_1(y)$, $s_2(y)$ and $s_4(y)$ are sufficient statistics for this analysis, and so $\pi(\mu, \sigma, \xi | s_{obs}) = \pi(\mu, \sigma, \xi | y_{obs})$ for these choices as there is no loss of information through the choice of summary statistics. However, both $s_1(y)$ and $s_2(y)$ are high-dimensional ($n = 49$), and so in practice, the resulting ABC posterior approximation may suffer from the curse of dimensionality. Both $s_3(y)$ and $s_4(y)$ are low dimensional, with the minimal one summary statistic per model parameter, but $s_3(y)$ is not sufficient for the model parameters, leading to some loss of information. Overall it might be expected that $s_4(y)$ will perform the best, as the dimension of the summary vector is minimised, without any loss of information (it is a minimal sufficient statistic). However, $s_4(y)$ will not typically be available in a typical ABC analysis due to the intractability of the likelihood function.

To complete the Bayesian model specification, the prior is notionally set as $p(\mu, \sigma, \xi) \propto 1$ over the support of the model parameters. To implement Algorithm 2 a sampling distribution, $g(\mu, \sigma, \xi)$, is required. For this analysis, the prior distribution is improper, and so setting $g(\mu, \sigma, \xi) = p(\mu, \sigma, \xi)$ is not appropriate. In these situations it is common to first identify a region of high posterior density based on a small number of initial simulations

of $(\theta^{(i)}, y^{(i)}) \sim L(y|\theta)b(\theta)$ where $b(\theta)$ is chosen freely by the user. Once the high posterior density region has been identified, then $g(\theta)$ can be specified to be e.g. uniform or proportional to the prior distribution over this region. See e.g. Fearnhead and Prangle (2012) for an explicit implementation of this idea. In the current setting, we exploit the fact that we know the actual location of the posterior distribution and specify $g(\mu, \sigma, \xi) = U(30, 70) \times U(5, 45) \times U(-0.3, 1.5)$.

The different scales and correlations of the summary statistics are taken into account by defining

$$\|s - s_{obs}\| = \left[(s - s_{obs})'\hat{\Sigma}^{-1}(s - s_{obs})\right]^{1/2}, \tag{9}$$

as the Mahalanobis distance, where $\hat{\Sigma}$ is an estimate of the covariance of s_{obs}. In practice the estimate of the covariance matrix only needs to be approximately correct, and so can be determined by identifying some point in parameter space (μ_0, σ_0, ξ_0) that is likely in an area of high posterior density, generating a number of summary statistic vectors $s^{(i)} \sim L(s|\mu_0, \sigma_0, \xi_0)$ conditional on this point, and then computing the sample covariance matrix of these vectors. See e.g. Luciani et al. (2009); Sisson and Fan (2011) for some examples of this approach. In the present analysis, (μ_0, σ_0, ξ_0) are set as the maximum likelihood estimates $(\hat{\mu}, \hat{\sigma}, \hat{\xi})$ for convenience.

The following results are based on $N = 1,000,000$ samples, where the smoothing kernel $K_h(u)$ is the uniform kernel over $(-h, h)$. The kernel scale parameter is determined as the 0.15, the 0.05 and the 0.005 quantile of $\|s_j^{(i)} - s_{j,obs}\|$ for each of the summary vectors $j = 1, \ldots, 4$, where $s_{j,obs} = s_j(y_{obs})$. This results in ABC approximations to the posterior distribution constructed from $150,000$, $50,000$ and $5,000$ samples respectively. Note that this retrospective definition of h is not quite in line with Algorithm 2, which strictly requires h to be determined before the first sample is drawn. However, this procedure of first drawing many $(\theta^{(i)}, s^{(i)})$ samples, and then determining h is commonplace in ABC practice (e.g. Beaumont et al. (2002); Blum et al. (2013)).

Figure 2 illustrates the estimated posterior marginal distributions of μ (left column), σ

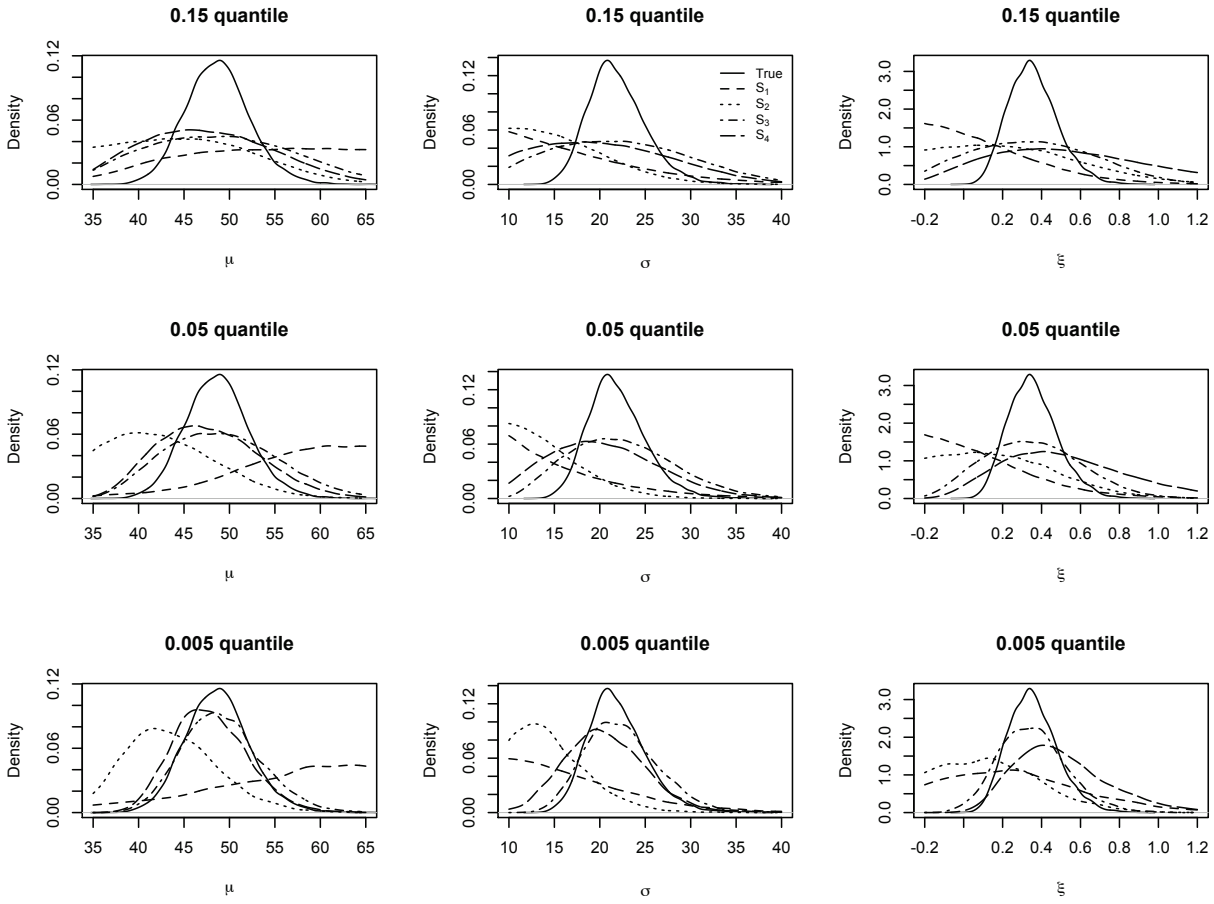

Figure 2: Estimates of the marginal posterior distributions of the GEV parameters μ (left column), σ (centre column) and ξ (right column) based on the Vargas dataset. Each panel shows the true posterior marginal distribution (solid line), and the four ABC approximations to the posterior based on the vectors of summary statistics $s_1(y), \ldots, s_4(y)$. Rows indicate the quantile of $\|s_j^{(i)} - s_{j,obs}\|$ used to determine the kernel scale parameter h.

(centre column) and ξ (right column) for both the true marginal posterior (solid line) and the four ABC marginal posterior approximations (various line types) based on $s_1(y), \ldots, s_4(y)$. Rows correspond to the different values of kernel scale parameter h, so that h is decreasing from the top row of Figure 2 to the bottom row.

It is readily apparent that all of the ABC approximations to the posterior distribution based on $s_1(y) = y$ (the dashed line) are very poor at best, regardless of the size of h. This should not be a surprise, as $\dim(s_1(y)) = 49$, and the chance of generating any 49-dimensional vector of independent observations, $s_1^{(i)}$, such that $\|s_1^{(i)} - s_{1,obs}\| \leq h$ becomes vanishingly small as h decreases. Of course, in theory, as $h \to 0$ then the resulting ABC posterior approximation will be equivalent to the true posterior distribution $\pi(\mu, \sigma, \xi|y_{obs})$ because $s_1(y)$ is a sufficient statistic. However, even taking the closest 5,000 samples (Figure 2, bottom row with the 0.005 quantile) out of 1 million was not enough to produce any reasonable accuracy in this case. This means that for practical purposes, it is not viable to use $s_1(y)$.

The ABC posterior approximations based on $s_2(y)$ (dotted lines) are as poor as those based on $s_1(y)$ for high values of h. However, there is some evidence that the density estimates are beginning to improve when h is reduced to the 0.005 quantile of $\|s_2^{(i)} - s_{2,obs}\|$. This is more obvious for the location parameter μ, which is typically the easiest parameter to estimate in any analysis, although there is little change for σ and ξ. The reason why the 49-dimensional vector $s_2(y)$ can achieve better performance than $s_1(y)$ (also in 49 dimensions) is that a close match of $s_2^{(i)}$ to $s_{2,obs}$ is far more likely to occur due to the induced dependence of the summary statistics. For example, consider that (say) the minimum values of each of y and y_{obs} are more likely to be close than a randomly chosen element from each dataset, which is the comparison being made when using $s_1(y)$. However, as with $s_1(y)$, the vector of summary statistics $s_2(y)$ still does not appear to be practically viable with a computational overhead of only $N = 1,000,000$ samples.

However, the vectors of 3 dimensional summary statistics $s_3(y)$ (dot-dash line) and $s_4(y)$

(long dashed line) both produce more practically useful ABC posterior approximations. The broad adherence of the approximate densities to the true posterior marginal distributions is apparent in Figure 2 even at the 0.15 quantile, and their accuracy only increases as h decreases. At the 0.005 quantile, either posterior is almost good enough to use in practice, although the L-moments based ABC approximation does appear to slightly better approximate the true posterior in this case, despite the information loss in $s_3(y)$. It seems quite likely that increasing N, thereby allowing h to be decreased further, would result in a viable ABC posterior approximation based on $s_3(y)$. While in theory using $s_4(y)$ should be slightly more efficient than using $s_3(y)$, this choice of summary statistics will not be available in practice.

In Section 2.3 a method is discussed – the regression adjustment – that will allow for an improved accuracy in an ABC analysis, without the need to increase computational overheads (N) further. For a technique to determine whether the chosen value of h is "low enough" for a good ABC posterior approximation when the true posterior distribution is unknown, see Prangle et al. (2014).

Overall, the main concepts of ABC have been illustrated by this simple analysis. ABC methods are themselves a simple and intuitive procedure that can produce viable approximations to posterior distribution without the need to evaluate the likelihood function. They work best when the vector of summary statistics are both highly informative for the model parameters, and are low dimensional. However, even then, ABC methods can have large computational overheads if high accuracy is desired. This is the price to pay for not being able to evaluate the likelihood function.

2.3 Other useful ABC methods

2.3.1 Regression-adjustment techniques

As previously discussed, standard ABC methods suffer from the curse of dimensionality (Blum, 2010), so that the kernel scale parameter h must increase for practical purposes as the number of summary statistics increases. As a result, there is still often a large discrepancy between $s^{(i)}$ and s_{obs} in the final sample from $\pi_h^{ABC}(\theta, s|s_{obs})$, whereas greatest accuracy of the ABC posterior approximation is obtained if $s^{(i)} \approx s_{obs}$ (i.e. if $h \to 0$). Regression-adjustment techniques aim to reduce this discrepancy by explicitly modelling the relationship between the sampled $\theta^{(i)}$ and $s^{(i)}$, and then adjusting $(\theta^{(i)}, s^{(i)}) \to (\theta^{(i)*}, s^{(i)*})$ so that $s^{(i)*} = s_{obs}$. The resulting (weighted) samples $\theta^{(i)*}$ will then form an improved approximation to $\pi(\theta|s_{obs})$ than those (weighted) $\theta^{(i)}$ from $\pi_h^{ABC}(\theta|s_{obs})$, if the regression model is correct.

The simplest form of a model for this is a homoscedastic regression in the region of s_{obs}, so that

$$\theta^{(i)} = m(s^{(i)}) + e^{(i)},$$

where $m(s^{(i)}) = \mathbb{E}[\theta|s = s^{(i)}]$ is the mean function, and the $e^{(i)}$ are zero mean, common variance random variates. In the case of a local-linear regression we have the model

$$m(s^{(i)}) = \alpha + \beta' s^{(i)}.$$

Beaumont et al. (2002) estimated this model by minimising the least squares criterion $\sum_{i=1}^{N} \omega^{(i)} \|m(s^{(i)}) - \theta^{(i)}\|^2$ where $\omega^{(i)} = K_h(\|s^{(i)} - s_{obs}\|)$. The regression-adjusted weighted sample $(\theta^{(i)*}, w^{(i)})$ drawn approximately from $\pi(\theta|s_{obs})$ is then obtained as

$$\theta^{(i)*} = \hat{m}(s_{obs}) + (\theta^{(i)} - \hat{m}(s^{(i)})), \tag{10}$$

where $\hat{m}(s) = \hat{\alpha} + \hat{\beta}' s$ is the fitted mean function. In addition to the local-linear regression-

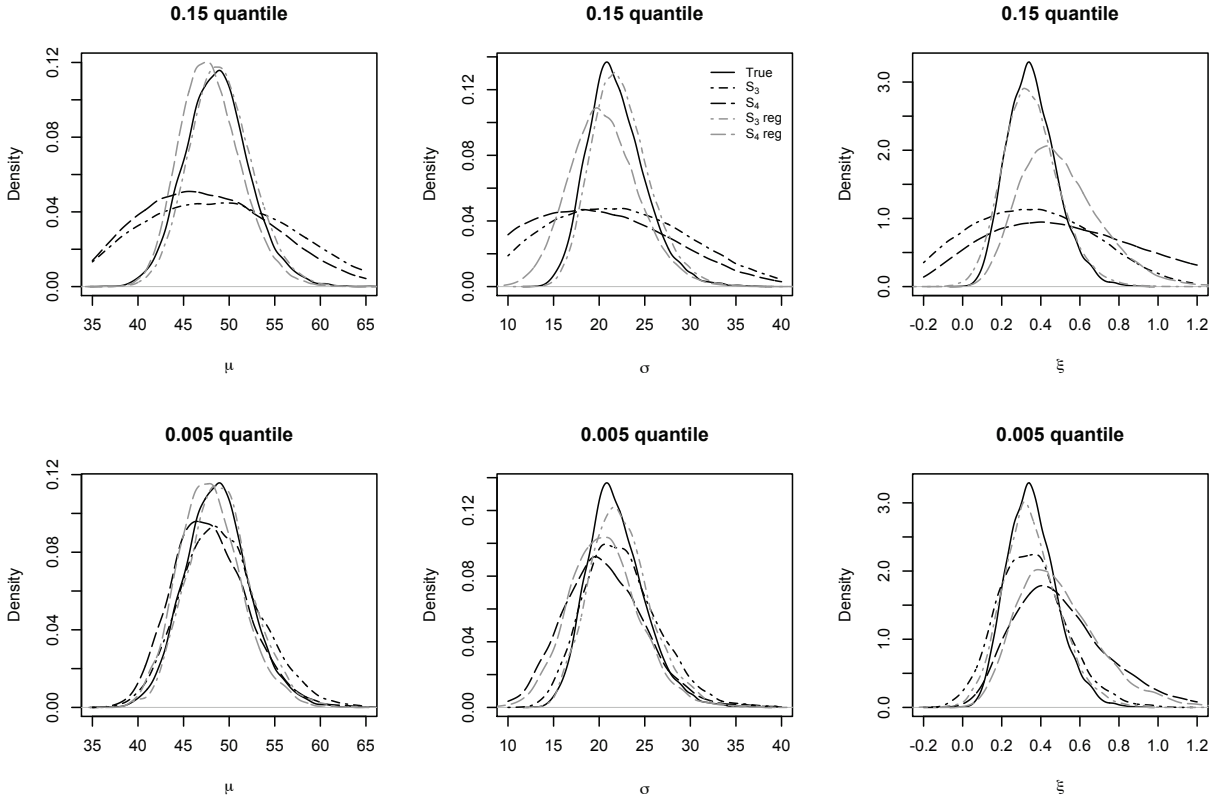

Figure 3: Regression-adjusted estimates of the marginal posterior distributions of the GEV parameters μ, σ and ξ based on the Vargas dataset. Each panel shows the true posterior marginal distribution (solid line) and the two ABC approximations to the posterior based on the vectors of summary statistics $s_3(y)$ and $s_4(y)$. ABC approximations without regression adjustment are illustrated with black lines. Those with a subsequent regression-adjustment are shown with grey lines.

adjustment of Beaumont et al. (2002), variations on this approach include a non-linear, heteroscedastic regression-adjustment (Blum and François, 2010) and a ridge regression-adjustment (Blum et al., 2013). A qualitatively different, but also useful form of marginal adjustment to improve ABC posterior approximation accuracy, that can be used in conjunction with the regression-adjustment is described by Nott et al. (2014).

Figure 3 illustrates the ABC approximations to the posterior marginal distributions of the GEV analysis described in Section 2.2, focusing on the 3-dimensional vectors of summary statistics $s_3(y)$ and $s_4(y)$. The black lines indicate the original approximations shown in Figure 2, whereas the grey versions of these lines show the same samples following the local-

linear regression adjustment (10).

In all cases, the regression-adjustment has made a noticeable improvement in the quality of the ABC posterior approximation, and in some cases the resulting density estimate is indistinguishable from the truth. This is a strong indicator that regression-adjustment methods should be routinely used in ABC analyses. It is interesting to note that the regression adjustment produces effectively the same approximation whether the initial samples $(\theta^{(i)}, s^{(i)})$ are obtained using relatively high (top row) or low (bottom row) values of kernel scale parameter h. This will occur when the fitted regression model remains effectively unchanged as h is varied. In this case, large computational savings can occur by obtaining an initial ABC approximation based on a larger value h, and then performing the adjustment, rather than doing the same but with a lower value of h.

2.3.2 Choice of summary statistics

Summary statistic identification is essentially a problem of dimension-reduction with respect to sufficiency for the target model. This is a huge research area in its own right. Blum et al. (2013) provide a recent comprehensive and comparative review of the methods that have been proposed for dimension reduction in ABC. These can be classified as best subset selection methods, projection techniques, regularisation approaches and other principled methods such as those based on indirect inference.

Some more recent developed approaches include those by Drovandi et al. (2014); Ruli et al. (2013) in addition to those particular cases where ABC can be reliably performed using the full original dataset, e.g. due to a likelihood factorisation (Barthelmé and Chopin, 2014; Bonassi et al., 2011; Fan et al., 2013; Jasra et al., 2012).

While improved techniques in this area are still being developed, it is probably accurate to say that there is no single "best" method. The optimum approach to identifying good summary statistics currently remains careful consideration of the specific model and analysis at hand.

2.3.3 Other ABC algorithms

The ABC method listed in Algorithm 2 is based on importance sampling. However there are many other standard algorithms for generating draws from posterior distributions, such as those based on MCMC and sequential Monte Carlo. Each of these approaches are beneficial in particular circumstances (Brooks et al., 2011; Doucet et al., 2001). A variety of ABC versions of these algorithms have been developed, and the basic mechanism by which they work is essentially the same as the importance sampling algorithm.

Essentially, each posterior simulation algorithm requires the computation of either importance weights, which are the ratio of the target distribution to the sampling distribution $\pi(\theta|y_{obs})/g(\theta)$, or MCMC acceptance probabilities which are constructed from the ratio of these, as

$$\pi(\theta'|y_{obs})g(\theta|\theta')/\left[\pi(\theta|y_{obs})g(\theta'|\theta)\right],$$

where the sampling distribution g is now a Markov proposal distribution. In the ABC framework, the approximate posterior distribution is augmented by the auxiliary datasets so that the target distribution is $\pi_h^{ABC}(\theta, s|s_{obs}) \propto K_h(\|s - s_{obs}\|)L(s|\theta)p(\theta)$ (7). This means that in any posterior simulation algorithm, the sampling or proposal distribution must also be defined on the same space, so that $q(\theta, s) \propto L(s|\theta)g(\theta)$ or $q(\theta, s|\theta', s') \propto L(s|\theta)g(\theta|\theta')$. As a result, when computing an importance weight or acceptance probability targeting $\pi_h^{ABC}(\theta, s|s_{obs})$, the intractable likelihood term $L(s|\theta)$ will always cancel out, as demonstrated at the end of Section 2. This leaves only computationally tractable terms, which can then be evaluated. Sisson and Fan (2011) demonstrate this in detail for the Metropolis-Hastings MCMC algorithm.

For further information on MCMC algorithms for ABC see e.g. Marjoram et al. (2003); Bortot et al. (2007); Sisson and Fan (2011), and for sequential Monte Carlo algorithms refer to e.g. Sisson et al. (2007); Beaumont et al. (2009); Drovandi and Pettitt (2011); Del Moral et al. (2012).

Alternative approaches to constructing an ABC approximation to the posterior distribution, based on building structured density estimates of the sample $(\theta^{(1)}, s^{(i)}), \ldots, (\theta^{(N)}, s^{(N)})$, have been developed by Bonassi et al. (2011) and Fan et al. (2013).

3 ABC for stereological extremes

In the production of clean steels, microscopically small particles termed *inclusions* are introduced during the production process. Metallurgic considerations indicate that the strength of the steel is directly related to the size of the largest inclusion in the block, and so inference on the largest inclusion size is important. Commonly, the sampling of inclusions involves measuring the maximum cross-sectional slice of each observed inclusion, $y_{obs} = (y_{1,obs}, \ldots, y_{n,obs})$, obtained from a two-dimensional planar slice through the steel block. Each cross-sectional inclusion size is greater than some measurement threshold, $y_{i,obs} > u$. The inferential problem is to analyse the unobserved distribution of the largest inclusion in the block, based on the information in the cross-sectional slice, y_{obs}. The focus on the size of the largest inclusion means that this is an extreme value variation on the standard stereological problem (Baddeley and Jensen, 2004).

Anderson and Coles (2002) proposed a mathematical model for those observed cross-sectional measurements. The proposed model assumed that the inclusions were spherical with diameters V, that their centres followed a homogeneous Poisson process with rate $\lambda > 0$ in the volume of steel, and that the inclusion diameters were mutually independent and independent of inclusion location. The distribution of the largest inclusion diameters, $V|V > v_0$, (i.e. those conditional on exceeding some threshold, v_0) was assumed to be well approximated by a generalised Pareto distribution, with distribution function

$$\Pr(V \leq v|V > v_0) = 1 - \left[1 + \frac{\xi(v - v_0)}{\sigma}\right]_+^{-1/\xi}, \tag{11}$$

for $v > v_0$, where $[a]_+ = \max\{0, a\}$, and $\sigma > 0$ and $-\infty < \xi < \infty$ are scale and shape parameters, following standard extreme value theory arguments (Coles, 2001). Accordingly the parameters of the full spherical inclusion model are $\theta = (\lambda, \sigma, \xi)$.

Each observed cross-sectional inclusion diameter, $y_{i,obs}$, is associated with an unobserved true inclusion diameter V_i. However, note that the probability of observing the cross-sectional diameter size $y_{i,obs}$ (where $y_{i,obs} \leq V_i$) is dependent on the value of V_i, as larger inclusion diameters give a greater chance that the inclusion will be observed in the two-dimensional planar cross-section. The number of observed inclusions, n, is also a random variable. In these terms, interest is in the distribution of the largest inclusion diameters, $V_1, \ldots, V_n$, given the observed cross-sectional measurements, $y_{1,obs}, \ldots, y_{n,obs}$.

Anderson and Coles (2002) were able to construct a likelihood function for this model by adapting the solution to Wicksell's corpuscle problem (Wicksell, 1925). They also overcome numerical difficulties with the resulting likelihood function of the model by treating the unobserved V_i as latent variables in a Bayesian hierarchical formulation. However, while their model assumptions of a Poisson process and inclusion independence are not unreasonable, the assumption that the inclusions are spherical is not plausible in practice.

Because of this, a generalisation of the spherical inclusion model was proposed by Bortot et al. (2007), who considered a family of ellipsoidal inclusions. In addition to the previous Poisson process and independence assumptions, this new model considered inclusions to be ellipsoidal and randomly oriented in space, with principal diameters (V^1, V^2, V^3). As before, the distribution function of $V^3 | V^3 > v_0$ is specified as the generalised Pareto distribution (11), where V^3 is defined as the largest ellipsoidal diameter. The other two principal diameters are defined as $V^j = U_j V^3$ for $j = 1, 2$, where U_1 and U_2 are independent uniform $U(0, 1)$ variables. Finally, in order to avoid ambiguity, the observed cross-sectional measurement $s_{i,obs}$ is assumed to be the largest principal diameter of the ellipse generated by the planar section of an ellipsoidal inclusion.

While the ellipsoidal inclusion is more realistic than the spherical inclusion model, there

are analytic and computational difficulties in extending likelihood-based inference to more general families of inclusion (Baddeley and Jensen, 2004; Bortot et al., 2007). As a result ABC methods are a good candidate procedure to approximate the posterior distribution $\pi(\lambda, \sigma, \xi | y_{obs})$ under the ellipsoidal model.

In the following analysis, the prior $p(\lambda, \sigma, \xi) \propto 1$ is specified as uniform over the support of the model parameters. For the purposes of implementing Algorithm 2 total of $N = 2$ million samples were drawn from the sampling distribution $g(\lambda, \sigma, \xi) = U(10, 80) \times U(0, 10) \times U(-3, 3)$ for the spherical model and $g(\lambda, \sigma, \xi) = U(60, 130) \times U(0, 10) \times U(-3, 3)$ for the ellipsoidal model based on a pilot analysis. A 7-dimensional vector of summary statistics was specified as

$$S(y) = (q_{0.5}(y), q_{0.7}(y), q_{0.9}(y), q_{0.95}(y), q_{0.99}(y), q_1(y), n') \tag{12}$$

where $q_a(y)$ denotes the a-th quantile of the dataset y (with interpolation if necessary), and n' is the number of observations in y. The Mahalanobis distance (9) is used to compare auxiliary and observed summary statistics $\|s - s_{obs}\|$, where the covariance matrix $\hat{\Sigma}$ is either the identity matrix I, or an estimate of $\mathrm{Cov}(s|\lambda_0, \sigma_0, \xi_0)$ based on 1,000 samples $s \sim L(s|\lambda_0, \sigma_0, \xi_0)$, where $(\lambda_0, \sigma_0, \xi_0) = (30, 1.5, -0.05)$ for the spherical model and $(\lambda_0, \sigma_0, \xi_0) = (95, 1.9, -0.1)$ for the ellipsoidal model. The values for $(\lambda_0, \sigma_0, \xi_0)$ were based on preliminary analyses for each model. The smoothing kernel $K_h(u)$ is the uniform kernel over $(-h, h)$ and the kernel scale parameter h is determined as the 0.001 quantile of the N samples $\|s^{(i)} - s_{obs}\|$, resulting in 2,000 samples from the approximate posterior $\pi_h^{ABC}(\lambda, \sigma, \xi | s_{obs})$. The observed dataset s_{obs} is derived from a set of 112 inclusion diameters previously analysed by Anderson and Coles (2002) and Bortot et al. (2007)

Figure 4 illustrates various marginal posterior estimates of $\pi(\lambda, \sigma, \xi | s_{obs})$ for both the spherical (top row) and ellipsoidal (bottom row) inclusion models. The ABC approximation to the posterior $\pi_h^{ABC}(\lambda, \sigma, \xi | s_{obs})$ using the identity matrix $\hat{\Sigma} = I$ is illustrated by the dashed

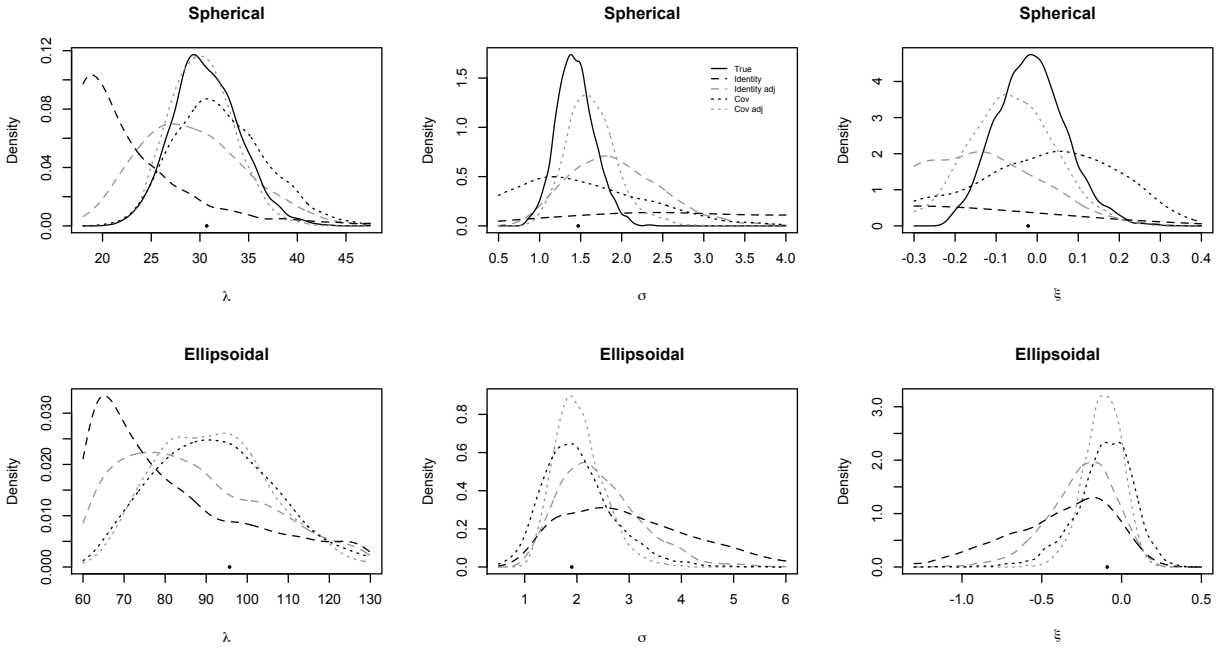

Figure 4: Marginal posterior estimates of λ, σ and ξ for the stereological extremes analysis. Top and bottom rows respectively correspond to the spherical and ellipsoidal inclusion models. Each panel shows the true posterior marginal distribution (solid line – spherical model only), and the ABC approximations to the posterior based on the summary statistics (12). Dashed lines indicate ABC posterior approximation using the identity matrix $\hat{\Sigma} = I$ in the Mahalanobis distance, dotted lines use an estimate of $\hat{\Sigma} = \mathrm{Cov}(s|\lambda_0, \sigma_0, \xi_0)$. Grey versions of each line indicate a subsequent regression-adjusted estimate. Points show the marginal posterior means estimated by Bortot et al. (2007).

lines, and the approximation using $\hat{\Sigma} = \mathrm{Cov}(s|\lambda_0, \sigma_0, \xi_0)$ is shown by the dotted lines. Grey lines denote the same approximations but with a subsequent regression adjustment. The solid black line indicates an estimate of the true posterior $\pi(\lambda, \sigma, \xi|y_{obs})$ following the latent variable model of Anderson and Coles (2002), for the spherical inclusion model only. The dots show the marginal posterior means estimated by Bortot et al. (2007).

For the spherical inclusions model (Figure 4, top row), there is a clear difference in the resulting estimates of the marginal posterior distributions due to the choice of covariance matrix in the comparison $\|s^{(i)} - s_{obs}\|$. When $\hat{\Sigma}$ approximates $\mathrm{Cov}(s|\lambda_0, \sigma_0, \xi_0)$, the ABC approximation (dotted lines) is able to locate the true posterior density (solid line) fairly well. However, when the identity matrix is used (dashed lines), the approximation is substantially worse. The primary reason for this is that the summary statistics with the highest variability – namely n' and the highest quantiles of y – dominate the comparison, so $\|s^{(i)} - s_{obs}\|$ is likely to be large unless these highly variable statistics happen to be close to those in s_{obs}. When summary statistic scaling (and correlation) is taken into consideration, the relative closeness of $s^{(i)}$ to s_{obs} can be better measured.

In addition, the implementation of a regression-adjustment also makes a substantial improvement in the quality of the ABC approximation, even for the very poor estimates using $\hat{\Sigma} = I$. The best results are clearly obtained when $\hat{\Sigma} \approx \mathrm{Cov}(s|\lambda_0, \sigma_0, \xi_0)$ is used, followed by a regression adjustment. Indeed, the inclusion rate parameter λ is almost perfectly estimated marginally. More detailed analysis (not shown) suggests that the results may improve further if the kernel scale parameter h could be lowered further. However, this would require a larger number of initial simulations, N.

For the ellipsoidal inclusions model (Figure 4, bottom row), while estimates of the true marginal posterior densities of $\pi(\lambda, \theta, \xi|y_{obs})$ are not available, qualitatively similar conclusions to the spherical inclusions model can be made. The best performing ABC approximation is likely to be when accounting for the scale and correlation of the summary statistics, followed by a regression adjustment. This approximation agrees well with the previous esti-

mates of marginal posterior means for these data given by Bortot et al. (2007)

A primary difference between the parameter estimates from the two models is that the number (rate) of inclusions is much higher for the ellipsoidal model. This reflects that the overall dimensions of an ellipsoid are smaller than a sphere with diameter the same as the largest principal diameter of the ellipsoid. Hence, given the observed planar intersections, a smaller rate of inclusion is predicted under the spherical inclusion model than under the ellipsoidal inclusion model. This analysis suggests that measures of inclusion impact that depend strongly on the rate of extreme inclusions are likely to be strongly sensitive to assumptions on inclusion shape.

While it is possible to perform Bayesian model selection through the computation of Bayes Factors in the ABC framework (Robert, 2011; Marin et al., 2014), in the present setting it is doubtful if the current data – the maximum cross-sectional slice of each observed inclusion – is informative for this quantity. To proceed further along the road of model choice, further measurements on each cross-sectional inclusion, such as the minimum and maximum diameters, would be required.

4 ABC and max-stable processes

4.1 Background

Suppose that there is interest in using ABC methods to fit max-stable processes to extreme temperature data in the context of actuarial risk estimation for a class of financial products known as weather derivatives. Weather derivatives are contracts that specify a weather reporting station, a time period, and payments corresponding to certain pre-determined weather events. The intention is for the buyer of this contract to be compensated by the seller if certain undesirable weather events occur (low snowfall for ski resorts, low rainfall for farms, high temperatures for electricity producers, etc.), but in practice, weather derivatives can also be bought and sold for purely speculative reasons.

Denoting the weather random variable by Y, then the resulting payment is also a random variable $P(Y)$. Here we are interested in derivatives with payments triggered by events $\{\max Y \geq u\}$ or $\{\min Y \leq u\}$ where Y is the variable of interest and u is some threshold meant to target extremes. For example, a payment P can be triggered if the maximum temperature in July exceeds 100 degrees Fahrenheit, with payment \$1000 for each degree in excess. This is shown mathematically as

$$P(Y) = \max\left(1000 \cdot \left(\max_{Y \in July}(Y) - 100\right), 0\right).$$

Because weather random variables Y can be positively correlated due to close proximity, financial outcomes of weather derivatives can also be positively correlated. When one considers the total payment from a collection of D weather derivatives $P = P_1 + \ldots + P_D$, recognition of the dependence is essential for fully characterising the distribution of P. The use of max-stable processes is one way to incorporate both spatial dependence and target extremes.

Max-stable processes are the infinite dimensional generalization of multivariate extreme value theory. Let $Z(x), x \in X \subseteq \mathbb{R}^p$ be a spatial process. If for all $n \geq 1$, there exists sequences $a_n(x), b_n(x)$, $x \in X$ such that for any $x_1, ..., x_D \in X$,

$$P^n\left(\frac{Z(x_d) - b_n(x_d)}{a_n(x_d)} \leq z(x_d), d = 1, ..., D\right) \to G_{x_1,...,x_D}(z(x_1), ..., z(x_D))$$

then $G_{x_1,...,x_D}$ is a multivariate extreme value distribution. If the above holds for all possible subsets $x_1, ..., x_D \in X$ for any $D \geq 1$, then the process is max-stable.

A method for constructing such processes was given by de Haan (de Haan, 1984; de Haan and Ferreira, 2006). Let $Y(x)$ be a non-negative stationary process on $\mathbb{R}^p$ such that $\mathbb{E}(Y(x)) = 1$ at each x. Let Π be a Poisson process on $\mathbb{R}_+$ with intensity $s^{-2}ds$. If $Y_i(x)$ are independent replicates of $Y(x)$, then

$$Z(x) = \max s_i \cdot Y_i(x), \quad x \in X$$

is a stationary max-stable process with unit Fréchet margins . Different choices of the process $Y(x)$ give different max-stable processes. Smith (1990) described a convenient interpretation of this process, with $\mathbb{R}^p$ as the space of storm centers, s_i as the magnitude of the i^{th} storm, and $Y_i(x)$ as the shape of the i^{th} storm. The maximum of independent storms is taken to be the max-stable process.

Schlather (2002) extended $Y(x)$ to be any stationary Gaussian process on $\mathbb{R}^p$ with correlation function $\rho(\cdot)$ and finite mean $\mu = \mathbb{E}[\max(0, Y(x))] \in (0, \infty)$. With s_i as a Poisson process on $(0, \infty)$ with intensity measure $\mu^{-1}s^{-2}ds$, the quantity

$$Z(x) = \max_i s_i \max(0, Y_i(x))$$

is a stationary max-stable process with unit-Fréchet margins. The bivariate distribution function is

$$P(Z_1 \leq z_1, Z_2 \leq z_2) = \exp\left[-\frac{1}{2}\left(\frac{1}{z_1} + \frac{1}{z_2}\right)\left(1 + \sqrt{1 - 2(\rho(h)+1)\frac{z_1 z_2}{(z_1+z_2)^2}}\right)\right] \tag{13}$$

where $\rho(h)$ is the correlation of the underlying Gaussian process Y and $h = ||x_1 - x_2||$. The correlation function may be chosen from one of the valid families of correlations for Gaussian processes, with one common choice being the Whittle-Matérn correlation,

$$\rho(h) = c_1 \frac{2^{1-\nu}}{\Gamma(\nu)}\left(\frac{h}{c_2}\right)^{\nu} K_\nu\left(\frac{h}{c_2}\right), \quad 0 \leq c_1 \leq 1, c_2 > 0, \nu > 0, \tag{14}$$

where c_1 is the nugget, c_2 is the range and ν is the smooth parameter. Figure 5 shows one realization of a process with the Whittle-Matérn correlation function.

To quickly summarize, we have a method for simulating realizations from max-stable processes, along with closed-form expressions for the bivariate distribution. Max-stable processes sampled at D locations form a $D-$dimensional MEVD, whose joint likelihood function can be expressed in terms of an exponent measure as in (3). If we were to evaluate

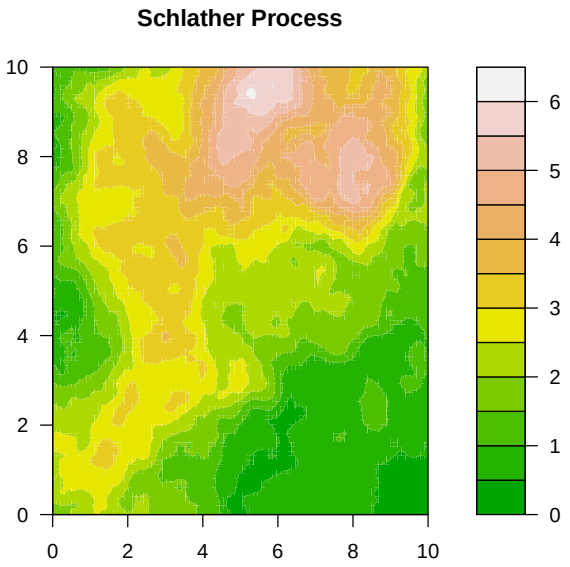

Figure 5: Extremal Gaussian process with Whittle-Matérn correlation with nugget $c_1 = 1$, range $c_2 = 3$, and smooth $\nu = 1$.

the probability that the process is below z at all locations $x_1, ..., x_D$, we get

$$P(Z(x_1) \leq z, ..., Z(x_D) \leq z) = \exp\left\{-\frac{\phi(x_1, ..., x_D)}{z}\right\},$$

where $\phi(x_1, ..., x_D) = V(1, ..., 1)$ is the *extremal coefficient* for the D locations. Since the bounds on the extremal coefficient $V(z_1, ..., z_D)$ are $1/z_1 + ... + 1/z_D$ and $\max(1/z_1, ..., 1/z_D)$, bounds on the extremal coefficient are D and 1, respectively, with a value of D corresponding to complete independence and a value of 1 corresponding to complete dependence. The value can be thought of as the number of effectively independent locations among the D under consideration. For a pair of locations, the extremal coefficient is

$$P(Z(x_1) \leq z, Z(x_2) \leq z) = \exp\left(-\frac{\phi(x_1, x_2)}{z}\right).$$

Observe that by using the bivariate distribution function available in closed form shown in

(13), one may write out the pairwise extremal coefficients explicitly as

$$\phi(h) \;=\; 1 + \left\{ \frac{1-\rho(h;\theta)}{2} \right\}^{1/2} \tag{15}$$

where $h = ||x_1 - x_2||$. Following Smith (1990) and Coles and Dixon (1999), if the field $Z(\cdot)$ has been transformed to unit-Fréchet, then $1/Z(\cdot)$ is unit exponential. This means $1/\max(Z(x_1), Z(x_2))$ is exponential with mean $1/\phi(x_1, x_2)$, and so a simple estimator of the extremal coefficient is

$$\hat{\phi}(x_1, x_2) = \frac{n}{\sum_{i=1}^{n} 1/\max(z_i(x_1), z_i(x_2))}$$

where i is the index for the block. Taking this idea further, Erhardt and Smith (2012) defined the tripletwise extremal coefficient

$$P(Z(x_j) \leq z, Z(x_k) \leq z, Z(x_l) \leq z) = \exp\left\{ -\frac{\phi(x_j, x_k, x_l)}{z} \right\}.$$

with estimator

$$\hat{\phi}(x_j, x_k, x_l) = \frac{n}{\sum_{i=1}^{n} 1/\max(z_i(x_j), z_i(x_k), z_i(x_l))}. \tag{16}$$

For D locations the number of tripletwise extremal coefficients is $\binom{D}{3}$, which grows quite rapidly as D increases. As part of an ABC analysis, lower dimensional summary statistics are preferred. Hence, a clustering step can be added to group these coefficients into a fixed number $K << \binom{D}{3}$ (see Erhardt and Smith (2012) for full details on this step). The $\binom{D}{3}$ triplet extremal coefficients may be estimated for the observed data using (16), and then these values are averaged within the K clusters. The result is the summary of the observed data, $s = (\bar{\phi}_1, ..., \bar{\phi}_K)$.

Following Algorithm 2, independent draws from the prior $\theta^{(i)} \sim p(\theta)$ are taken. For each draw from the prior, a max-stable process $Z^{(i)}$ with unit-Fréchet margins is simulated on the same locations and for the same number of years as the observed data. For this data $Z^{(i)}$, all tripletwise extremal coefficients are estimated, and averages within each cluster group are

taken to produce $s^{(i)} = (\bar{\phi}_1, ..., \bar{\phi}_K)$. To compare summaries, the sum of absolute deviations is used a distance measure, so that

$$\|s^{(i)} - s_{obs}\| = \sum_{k=1}^{K} |s_k^{(i)} - s_{k,obs}|$$

where $s_k^{(i)}$ and $s_{k,obs}$ denote the k-th element of $s^{(i)}$ and s_{obs} respectively. As before, the smoothing kernel $K_h(u)$ is uniform over $(-h, h)$, and the kernel scale parameter h is determined as the 0.01 quantile of the N samples $\|s^{(i)} - s_{obs}\|$.

4.2 Application

The observed data are annual maxima daily temperatures taken from 39 locations in the midwestern United States with complete summer (June 1 - August 31) temperature records from 1895 to 2009. All sites are located between longitudes 93 and 103 degrees west, and latitudes 37 to 45 degrees north.

Each of the 39 marginal distributions were transformed to unit-Fréchet through ordinary maximum likelihood estimation. Based on the dependence model (14) with no nugget effect ($c_1 = 1$), the prior for the spatial dependence parameter vector $\theta = (c_2, \nu)$ is specified as $[0, 7] \times [0, 7]$. This choice aims to represent vague non-informativeness and also place positive support on a large range of possible parameter values. Based on candidate draws $\theta^{(i)} \sim p(\theta)$, $i = 1, \ldots, N$, max-stable processes $Z^{(i)} \sim f(Z|\theta^{(i)})$ were simulated for 115 years at the same 39 locations.

The ABC posterior approximation $\pi_h^{ABC}(\theta|y_{obs})$ based on $N = 100,000$ samples is shown in Figure 6. The approximate posterior has correctly identified the credible regions of the parameter space from these data, retaining only draws from a crescent shaped subspace. Inversely changing values of c_2 and ν can produce similarly shaped correlation functions $\rho(h)$ – hence the crescent shape for the posterior approximation. The posterior predictive distribution of spatial correlation functions $\rho(h)$ is illustrated in the lower panel of Figure 6.

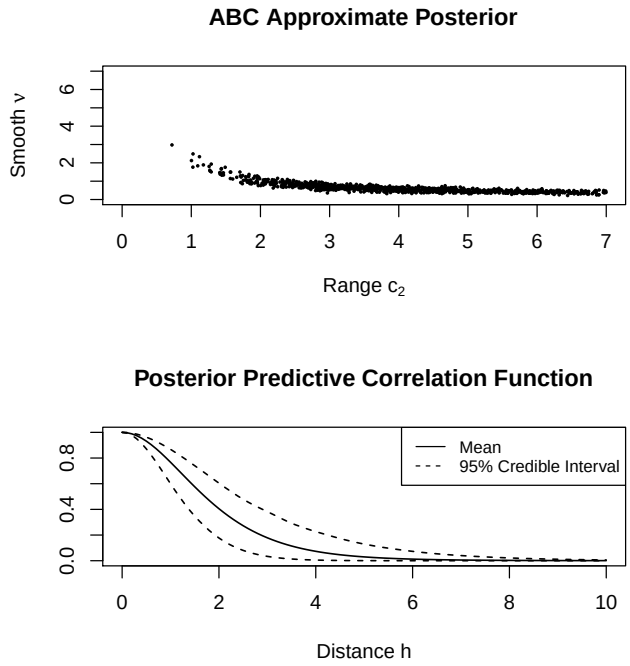

Figure 6: Top: Samples drawn from the ABC approximate posterior distribution based on fitting a max-stable process. The uniform prior for (c_2, ν) is the full range displayed on the scatterplot. Bottom: Posterior predictive distribution of the corresponding spatial correlation function $\rho(h)$ following equation (14). The solid line shows the pointwise posterior mean, and the dashed lines show the pointwise 2.5% and 97.5% quantiles.

This approximate posterior may now be used to fully incorporate the parameter uncertainty of θ into estimates of the distribution of total loss $P = P_1 + \ldots + P_D$ and associated actuarial risk measures.

5 Discussion

This chapter has outlined how approximate Bayesian statistical modeling of extremes can be viable in situations where the likelihood function either cannot be evaluated numerically, or written in closed form. Sections 3 and 4 illustrate ABC methods as applied to two different models in extremes, with various model and data complexities, in which an intractable likelihood is a common issue. All that ABC methods require in order to be implemented is a fast algorithm to simulate realisations from the intractable likelihood, and a suitable

vector of summary statistics. The former requirement is not always a trivial procedure for extremal models, and commonly needs some subtle theory and computational tricks. However, procedures for accurately simulating observations from extremal models continues to be an area of active research.

Perhaps the most important aspect of any ABC analysis is the identification of a low-dimensional vector of summary statistics that are highly informative for the model parameters. Once the summary statistic is chosen, the best possible ABC approximation to the posterior is $\pi(\theta|s_{obs}) \approx \pi(\theta|y_{obs})$, and this occurs when the kernel scale parameter $h \to 0$. If the summary statistic is poorly chosen, the rest of the ABC procedure (e.g. Algorithm 2) can not recover the loss of information. In fact, as $h > 0$ is a necessity in practice, the best ABC posterior approximation achievable is $\pi_h^{ABC}(\theta|s_{obs})$ which is in general a worse approximation than $\pi(\theta|s_{obs})$. Hence, the choice of summary statistics is critical. To date, probably the most useful and principled general approach to identify summary statistics is developed by Fearnhead and Prangle (2012), whereas a method to derive summary statistics from the score function of a related composite model (such as a pairwise likelihood) could be of particular interest for modelling spatial extremes (Ruli et al., 2013).

Coupled with the choice of summary statistics is the issue of computational overheads. ABC methods are a classic example of a computation versus accuracy trade-off. This implies that any developments in improving the efficiency of an ABC analysis can be converted into producing a more accurate inference for the same computational cost. Some more sophisticated algorithms than importance sampling are detailed in Section 2.3.3, but other ideas, such as terminating the generation of a dataset early if it is likely to be rejected (Prangle, 2014) are also starting to be developed. Improving the "approximation" (in Approximate Bayesian Computation) while reducing the "computation" is an active research area in ABC.

It is worth stating that the development of a tractable, closed-form joint likelihood for a stochastic process would entirely obviate the need for ABC methods by allowing for a more conventional Bayesian analysis based on direct usage of the likelihood. A flavour of

this evolution can be found in both Genton et al. (2011) and Aandahl et al. (2014). After a period of roughly two decades in which researchers could write down only the bivariate joint distribution function for the Gaussian max-stable process, Genton et al. (2011) extended the closed-form to include the trivariate joint likelihood. In the setting of population genetics, Aandahl et al. (2014) show that models that were previously only available for analysis using ABC, have subsequently become tractable with further analytical development. These are useful reminders that what is considered intractable today may not remain intractable in the future. In the meantime, however, the statistician's toolbox remains empowered through the availability of ABC methods.

References

Aandahl, R. Z., T. Stadler, S. A. Sisson, and M. M. Tanaka (2014). Exact vs. approximation computation: Reconciling different estimates of mycobacterium tuberculosis epidemiological parameters. *Genetics 196*, 1227–1230.

Anderson, C. W. and S. G. Coles (2002). The largest inclusions in a piece of steel. *Extremes 5*, 237–252.

Baddeley, A. and E. B. V. Jensen (2004). *Stereology for Statisticians*. Chapman & Hall/CRC.

Barthelmé, S. and N. Chopin (2014). Expectation-propagation for likelihood-free inference. *Journal of the American Statistical Association 109*, 315–333.

Beaumont, M. A. (2010). Approximate Bayesian computation in evolution and ecology. *Annu. Rev. Ecol. Evol. System. 41*(1).

Beaumont, M. A., J.-M. Marin, J. M. Corunet, and C. P. Robert (2009). Adaptivity for ABC algorithms: The ABC-PMC scheme. *Biometrika 96*, 983–990.

Beaumont, M. A., W. Zhang, and D. J. Balding (2002). Approximate Bayesian computation in population genetics. *Genetics 162*, 2025–2035.

Bertorelle, G., A. Benazzo, and S. Mona (2010). ABC as a flexible framework to estimate demography over space and time: some cons, many pros. *Molecular Ecology 19*(13), 2609–2625.

Blum, M. G. B. (2010). Approximate Bayesian computation: A nonparametric perspective. *Journal of the Amercan Statistical Association 105*, 1178–1187.

Blum, M. G. B. and O. François (2010). Non-linear regression models for approximate Bayesian computation. *Statistics and Computing 20*, 63–73.

Blum, M. G. B., M. A. Nunes, D. Prangle, and S. A. Sisson (2013). A comparative review of dimension reduction methods in approximate Bayesian computation. *Statistical Science 28*(189-208).

Bonassi, F. V., L. You, and M. West (2011). Bayesian learning from marginal data in bionetwork models. *Statistical Applcations in Genetics and Molecular Biology 10*, Article 10.

Bortot, P., S. G. Coles, and S. A. Sisson (2007). Inference for stereological extremes. *Journal of the Amercan Statistical Association 102*, 84–92.

Brooks, S. P., A. Gelman, G. L. Jones, and X.-L. Meng (Eds.) (2011). *Handbook of Markov Chain Monte Carlo*. CRC Press.

Chen, M.-H., Q.-M. Shao, and J. G. Ibrahim (2000). *Monte Carlo Methods in Bayesian Computation*. Springer.

Coles, S., J. Bawa, L. Trenner, and P. Dorazio (2001). *An introduction to statistical modeling of extreme values*, Volume 208. Springer.

Coles, S. G. (2001). *An introduction to statistical modeling of extreme values*. Springer.

Coles, S. G. and M. J. Dixon (1999). Likelihood-based inference for extreme value models. *Extremes 2*(1), 5–23.

Coles, S. G. and L. R. Pericchi (2003). Anticipating catastrophes through extreme value modelling. *Journal of the Royal Statistical Society, Series C 52*, 405–416.

Coles, S. G., L. R. Pericchi, and S. A. Sisson (2003). A fully probabilistic approach to extreme value modelling. *Journal of Hydrology 273*, 35–50.

Coles, S. G. and E. A. Powell (1996). Bayesian methods in extreme value modelling: a review and new developments. *International Statistical Review/Revue Internationale de Statistique*, 119–136.

Csilléry, K., M. G. B. Blum, O. E. Gaggiotti, and O. François (2010). Approximate Bayesian computation (ABC) in practice. *Trends Ecol. Evolut. 25*(7), 410–418.

de Haan, L. (1984). A spectral representation for max-stable processes. *The Annals of Probability*, 1194–1204.

de Haan, L. and A. Ferreira (2006). Extreme value theory: An introduction. 2006.

Del Moral, P., A. Doucet, and A. Jasra (2012). An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Statistics and Computing 22*, 1009–1020.

Doucet, A., N. de Freitas, and N. Gordon (Eds.) (2001). *Sequential Monte Carlo Methods in Practice*. Springer.

Drovandi, C. C. and A. N. Pettitt (2011). Likelihood-free Bayesian estimation of multivariate quantile distributions. *Computational Statistics & Data Analysis 55*, 2541–2556.

Drovandi, C. C., A. N. Pettitt, and A. Lee (2014). Bayesian indirect inference using a parametric auxiliary model. *Statistical Science*, in press.

Erhardt, R. J. and R. L. Smith (2012). Approximate bayesian computing for spatial extremes. *Computational Statistics & Data Analysis 56*(6), 1468–1481.

Fan, Y., D. J. Nott, and S. A. Sisson (2013). Approximate Bayesian computation via regression density estimation. *STAT 2*, 34–48.

Fearnhead, P. and D. Prangle (2012). Constructing summary statistics for approximate bayesian computation: semi-automatic approximate bayesian computation. *Journal of the Royal Statistical Society, Series B 74*, 419–474.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin (Eds.) (2003). *Bayesian Data Analysis* (2nd ed.). Chapman & Hall/CRC.

Genton, M. G., Y. Ma, and H. Sang (2011). On the likelihood function of gaussian max-stable processes. *Biometrika 98*(2), 481–488.

Gnedenko, B. V. (1944). Limit theorems for sums of independent random variables. *Uspekhi Matematicheskikh Nauk* (10), 115–165.

Hosking, J. R. M. (1990). L-moments: Analysis and estimation of distributions using linear combinations of order statistics. *Journal of the Royal Statistical Society, Series B 52*, 105–124.

Jasra, A., S. S. Singh, J. S. Martin, and E. McCoy (2012). Filtering via approximate Bayesian computation. *Statistics and Computing 22*, 122–1237.

Leadbetter, M. and G. Lindgren (1983). Rootz en, h.(1983). extremes and related properties of random sequences and processes.

Luciani, F., S. A. Sisson, H. Jiang, A. R. Francis, and M. M. Tanaka (2009). The epidemiological fitness cost of drug resistance in *mycobacterium tuberculosis*. *Proceedings of the National Academy of Sciences of the USA 106*, 14711–14715.

Marin, J.-M., N. Pillai, C. P. Robert, and J. Rousseau (2014). Relevant statistics for Bayesian model choice. *Journal of the Royal Statistical Society, Series B*, in press.

Marin, J.-M., P. Pudlo, C. P. Robert, and R. J. Ryder (2012). Approximate Bayesian computational methods. *Statist. Comput. 22*, 1167–1180.

Marjoram, P., J. Molitor, V. Plagnol, and S. Tavaré (2003). Markov chain Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences of the USA 100*, 15324–15328.

Nott, D. J., Y. Fan, L. Marshall, and S. A. Sisson (2014). Approximate Bayesian computation and Bayes linear analysis: Towards high-dimensional abc. *Journal of Computational and Graphical Statistics 23*, 65–86.

O'Hagan, A. and J. Forster (2004). *Bayesian Inference* (2nd ed.), Volume 2B of *Kendall's Advanced Theory of Statistics*. Arnold.

Pickands, J. (1981). Multivariate extreme value distributions. In *Proceedings 43rd Session International Statistical Institute*, Volume 2, pp. 859–878.

Prangle, D. (2014). Lazy ABC. *http://arxiv.org/abs/1405.7867*.

Prangle, D., M. G. B. Blum, G. Popovic, and S. A. Sisson (2014). Diagnostic tools for approximate Bayesian computation using the coverage property. *Australia and New Zealand Journal of Statistics*, in press.

Pritchard, J. K., M. T. Seielstad, A. Perez-Lezaun, and M. W. Feldman (1999). Population growth of human Y chromosomes: A study of Y chromosome microsatellites. *Molecular Biology Evolution 16*, 1791–1798.

Reeves, R. and A. N. Pettitt (2005). A theoretical framework for approximate Bayesian computation. In A. R. Francis, K. M. Matawie, A. Oshlack, and G. K. Smythe (Eds.), *Statistical Solutions to Modern Problems: Proceedings of the 20th International Workshop on Statistical Modelling, Sydney, 10-15 July, 2005*. University of Western Sydney Press, Sydney.

Resnick, S. (1987). Extreme values, regular variation, and point processes (1987).

Robert, C. P. (2011). Lack of confidence in approximate Bayesian computation model choice. *Proceedings of the National Academy of Sciences of the USA 108*, 15112–15117.

Ruli, E., N. Sartori, and L. Ventura (2013). Approximate Bayesian computation with composite score functions. *www.arxiv.org/abs/1311.7286*.

Schlather, M. (2002). Models for stationary max-stable random fields. *Extremes 5*(1), 33–44.

Sisson, S. A. and Y. Fan (2011). Likelihood-free Markov chain Monte Carlo. In S. P. Brooks, A. Gelman, G. Jones, and X.-L. Meng (Eds.), *Handbook of Markov Chain Monte Carlo*, pp. 319–341. Chapman and Hall/CRC Press.

Sisson, S. A., Y. Fan, and M. M. Tanaka (2007). Sequential Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences of the USA 104*, 1760–1765. Errata (2009), 106, 16889.

Smith, R. L. (1990). Max-stable processes and spatial extremes. *Unpublished manuscript, Univer*.

Tavaré, S., D. J. Balding, R. C. Griffiths, and P. Donnelley (1991). Infering coalescent times from DNA sequence data. *Genetics 145*, 505–518.

Wicksell, S. D. (1925). The corpsucle problem: A mathematical study of a biometric problem. *Biometrika 17*, 84–99.

Wieczorek, G. F., M. C. Larsen, L. S. Eaton, B. A. Morgan, and J. L. Blair (2001). Debris-flow and flooding hazards associated with the December 1999 storm in coastal Venezuela and strategies for mitigation. *U. S. Geological Survey*, pubs.usgs.gov/of/2001/ofr–01–0144/.

Wilkinson, R. D. (2008). Approximate Bayesian computation (ABC) gives exact results under the assumption of model error. Technical Report arXiv:0811.3355v1, Arxiv.